Revista Colombiana de Computación
Vol. 25, No. 1. January – June 2024, pp. 19-28
e-ISSN: 2539-2115, https://doi.org/10.29375/25392115.5058
Selected paper previously presented at the Latin America High Performance Computing Conference (CARLA 2023), an event held in Cartagena de Indias, Colombia, September 18-22, 2023.

Analysis Model of Energy Consumption Variables for Data Processing in High-Performance Computing Systems

Jorge Lozoya Arandia[1], Carlos Jesahel Vega Gómez[2], Lester Antonio Acevedo Montoya[3] and Verónica Lizette Robles Dueñas[4]

[1] School of Engineering and Technological Innovation, University of Guadalajara, CUSUR, Ciudad Guzman, Jalisco 45425, Mexico,
[2] School of Engineering and Technological Innovation, University of Guadalajara, CUTLAJO, Tlajomulco, Jalisco 45425, Mexico,
[3] School of Engineering and Technological Innovation, University of Guadalajara, CUTonala, Tonala, Jalisco 45425, Mexico,
[4] Data Analysis and Supercomputer Center, University of Guadalajara, Zapopan Jalisco, Mexico
jorge.larandia@academicos.udg.mx, jesahel.vega@academicos.udg.mx, lester.acevedo9845@academicos.udg.mx, lizette.robles@udg.mx

(Received: 6 February 2024; accepted: 18 June 2024, Published online: 30 June 2024)

Abstract. One of the main challenges in the efficient operation of a high-performance computing (HPC) center is the energy consumption generated by the operation of the data center where the HPC equipment is housed, mainly because this consumption is reflected in very high accounts payable, and this may affect the level of service offered to users. The study of the different factors and elements that can make energy consumption more efficient in these data centers provides an opportunity to focus these resources on elements that favor the use of HPC. The design variables provided by manufacturers to manage HPC systems and monitoring systems provide an accurate view of the behavior of these variables according to how they are used. HPC architectures are configured in a very particular way for each HPC data center, creating particular scenarios of operation and performance in each implementation. Various proposals and technologies have been developed for the analysis of the energy consumption of a data center, and the processing elements include a series of indicators and technologies that manufacturers have developed to determine the energy efficiency. This article seeks to identify this series of processing and performance variables, which affect the energy consumption of HPC equipment, for the implemented computing architectures based on the analysis of performance models to obtain a general over-view of their effect on energy consumption in a case study to identify the behaviors of particular job assignment factors and provide an analysis of the energy consumption under particular conditions.

Keywords: High-performance Computing, Data Center, Energy Simulation, Power Usage Effectiveness.

1 Introduction

Data centers that house high-performance computing (HPC) equipment have very particular characteristics when it comes to their computing architectures, data networks and cooling systems, as well as the management of their services. These architectures are based on the very special characteristics of each implementation, which generate very different and complex implementations for each piece of equipment; these architectures are made up of hardware and software for operating and monitoring the performance of each element, which is why standardized performance analysis models must be individualized to each piece of equipment implemented. Coupled with this, the speed at which both hardware and software updates emerge generates a continuous update dynamic: "The appearance of new Instruction Set Architectures (ISAs) in most recent High-Performance Computing (HPC) systems, together with the layers of system

©2024 Universidad Autónoma de Bucaramanga - UNAB. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY-NC-SA 4.0) license (https://creativecommons.org/licenses/by-nc-sa/4.0/).

software and complex scientific applications running on top of them, makes the performance and power figures challenging to evaluate" (Criado, et al., 2020).

Due to these new computing architectures in HPC and different processes and software, it is important to implement precise methodologies to make the most of the installed computing capacities; the use of optimization and performance improvement methodologies make it possible to analyze all these variables together and to interpret the behavior of each of the elements of an HPC system. This is achieved using methodologies such as those proposed by the EU Center of Excellence for Performance Optimization and Productivity (POP) (Wagner, Mohr, Giménez, & Labarta, 2019). Generating a structured analysis of these variables by applying methodologies such as those mentioned above to the metrics collected from an implemented system provides practical results concerning these models. The performance of simultaneous multi-threading (SMT) strongly depends on the architecture and the compiler (Banchelli, Garcia-Gasulla, Houzeaux, & Mantovani, 2020). These variables and monitoring systems make it possible to analyze the performance and energy consumption of the elements of the installed HPC system.

This article is focused on modeling the behaviors that affect the energy consumption of the HPC equipment in order to carry out an analysis of this energy consumption and the variables that affect it. The energy consumption of HPC data centers is one of the most important issues in the operation of these centers (D'Agostino, et al., 2019) and is a factor of great importance when it comes to evaluating the efficiency and productivity of the equipment housed in a data center and the data processing that is carried out. The analysis of the energy consumption of a data center has to be focused directly on the equipment that is housed in it, and in the case of a specialized data center such as an HPC data center, it must also be known how the processing and storage operations and the use of the data network are carried out. It is evident that different resource optimization techniques must be selected for underloaded and overloaded hosts depending on the servers and user data type (Hijji, et al., 2022). This approach must be applied to appropriate optimization techniques to obtain a result for the analyzed architecture.

Processing in HPC systems involves performing different processing jobs that are made up of a series of executions and algorithms according to the user's needs; these jobs are assigned a number of central processing units (CPUs), and the CPUs are grouped by nodes (see Fig. 1). These jobs are run on different software platforms and have different execution times. Due to the fact that the jobs being executed have different performances and different pieces of software that support the jobs on the HPC processing equipment, analyzing them requires that an understanding of all the processing teams (Funika, Zientarski, Badia, Labarta, & Bubak, 2008). These characteristics generate different energy consumption behaviors and are analyzed using the same behavior variables that are generated by performance control systems.

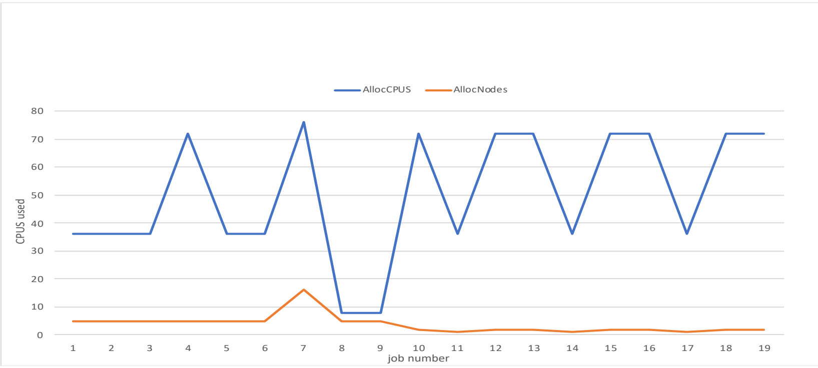

Figure 1. CPUs and nodes used for each job assigned in an HPC system.

Performing frequent analyses of the behaviors and performance of the HPC equipment that is being operated makes it possible to visualize and prevent behaviors that may derive from a physical or configuration failure: "Performance Audits provide an initial analysis and overview that measures a range of performance metrics to assess quality of performance and identify the issues affecting performance" (Wagner, Mohr, Giménez, & Labarta, 2019).

An important tool for maintaining efficient energy consumption is the use of performance audits, since they allow one to visualize the energy consumption and the behaviors that cause increased energy consumption. The performance of a job de-pends on how well these functions can be approximated; the first dependence is im-posed by the specific application, and the second is related to the physical model and basis set employed (Mohr, et al., 2017; Dawson, Mohr, Ratcliff, Nakajima, & Genovese, 2020).

The analysis of performance variables in HPC equipment generates a series of energy consumption behaviors related to processing actions and job assignment. In this way, an analysis of the behavior of the processor and its control actions can be generated. The assignment of processes and the different characteristics of each HPC architecture and the type of software and specific driver for each job should be considered. The analysis of these architectures can be combined with different techniques to develop projections of the energy consumption, which can be used to optimize the energy consumption. Deep learning techniques have been used to under-stand the balance of HPC infrastructure in terms of speed and energy consumption (Carastan-Santos & Pham, 2022).

For this article, a series of performance analysis and monitoring variables are pro-posed for the purpose of identifying how they interact with each other and influence the energy consumption of HPC equipment, so that the energy consumption can be characterized for analysis and the relationships between these variables and the different actions that make up HPC processing can be determined.

1.1 Center for Data Analysis and Supercomputing at the University of Guadalajara

The Data Analysis and Supercomputing Center (CADS) of the University of Guadalajara is a high-performance research center serving the scientific community and has HPC infrastructure housed in a specialized data center that allows the proper operation of the equipment (CADS, 2018). The installed computing architecture consists of a series of parallel Xeon Gold processors, Nvidia Tesla P100 processors and other Xeon Phi equipment (Table 1).

The allocation of computing resources to each job is done centrally and corresponds to the particular requirements of each project, which refer to the processing, storage and memory requirements. This administration is carried out through process control and scheduling. This system allows the monitoring of the operating parameters of the infrastructure and the energy consumption associated with the different systems.

For this article, the energy consumption was projected and the behavior of the HPC system installed in CADS was analyzed; specifically, different jobs with different resource allocations and different software components and functions were analyzed.

Table 1. Description of CADS infrastructure processors.

Processor	Cores	Memory
Xeon-6154 (Skylake)	18 cores at 3.0 GHz	192–392 GB RAM
Xeon-6154 (Skylake)	18 cores at 3.0 GHz	512 GB RAM
Intel Xeon Phi 7250		192 GB DDR4 16 GB MCDRAM
Xeon Gold-6154 (Skylake)	18 cores at 3.0 GHz	64 GB DDR4 PCI Nvidia Tesla P100
Intel Xeon E5-2650v4 12C/24T 2.20 GHz		128 of memory

2 Energy Performance Analysis in HPC Processing

This article proposes a model structured using stages based on specialized performance analysis software associated with user jobs, each of which has special characteristics. The first stage involves taking measurements using the sensors installed in the HPC monitoring system of the CADS at the University of Guadalajara; the measurements of the variables provided by the system were characterized according to the assigned jobs. These variables are defined in a particular way, so that once the crossovers are made, we can identify what causes energy consumption. Once this characterization is complete, overall and focused analyses are carried out.

The goal is to identify behaviors and relate them directly to energy consumption. The variables are identified and a comparative performance modeling table is generated, projecting the impact of each job according to its defined characteristics and allowing a detailed projection analysis. The platforms generate reports concerning these variables, and they will be analyzed to identify proposed solutions and conclusions.

The analysis of the energy consumption is vital to identifying how to achieve efficiency in HPC systems. This has become a high priority due to their complexity: "Energy efficiency is already a major concern in the design of computer systems, especially in the design of Exascale systems. There are a few activities targeting his problem" (Jarus, Varrette, Oleksiak, & Bouvry, 2013).

The measurement survey is developed using the Slurm monitoring program by registering the equipment housed in the data center; likewise, the energy consumption of the load centers connected to the same equipment is monitored. The work structure and use of the cluster are based on nodes, which are logically ordered according to an assigned number of processors. Once the nodes are structured, the writing and reading processing behaviors associated with each job must be monitored, and the number of CPUs used by each node is identified.

The workload description can be used to obtain the relationship between the function performed by the processor, the time it takes to perform this function and the energy consumption that occurs while this function is being performed; these data allow us to associate different architectures and the software being used with the hardware function and to determine how it affects the energy consumption (Table 2). The time used by each process is directly related to its energy consumption and allows the identification of behaviors that are not consistent with the processor design or the type of job

Table 2. Work load variables description.

Function	Description	Energy consumption Relationship
AllocCPUS	Total number of CPUs allocated to the job	Energy per CPU
AllocNodes	Total number of nodes allocated to the job	Total energy per node
CPUTimeRAW	Time used (Elapsed time * CPU count) by a job or step in cpu-seconds	Time of use
Avedisc read	Average disk usage read	Energy associated used
Avedisk write	Average disk usage write	Energy associated used
ConsumedEnergy	Consumed energy per job	Energy consumption
CPUTime	Tiem of cpu usage	Time per CPU

The analysis of different studies indicates that the performance of HPC equipment is important, and its various architectures can be affected by the elements that compose it: "The benchmarking results demonstrate the importance of sophisticated memory allocation on machines" (Plesser, Eppler, Morrison, Diesmann, & Gewaltig, 2007). In the case of memory use, using the disk writing and disk reading variables, we can identify correlations and customize the energy consumption.

Monitoring the HPC process causes a large amount of disk reading and writing, as well as a large amount of CPU usage assigned to the nodes, although this operation is complex; in some studies, it has been shown that this operation causes fewer problems. Other studies have achieved good scaling for large-scale simulations on systems with thousands of processors, albeit on less difficult problems (Djurfeldt, et al., 2005; Migliore, Cannia, Lytton, Markram, & Hines, 2006). This should generate energy consumption behaviors of the same order. The different architectures allow this energy consumption to be made more efficient using the process controllers and the workloads to be executed, and "even in kernels where the CPU reaches better runtimes, the FPGA counterpart is more energy efficient" (Favaro, Dufrechou, & Oliver, 2022). The analysis of the set of factors of the hardware and software in an HPC system generates different behavior conditions and variations in the energy consumption that respond to these particularities.

3 HPC Workload Energy Consumption

Given the complexity of the factors to be analyzed in an HPC system, it is important to identify the processes to be analyzed and then to define the variables of their behavior; the different monitoring tools become the data collection tools used to measure these variables. This article analyzes these variables for each assigned job in order to identify software and hardware factors in processing.

In this work, we only analyze the processing system; however, studies related to energy consumption in HPC systems have found that another of the factors of great importance is the use of the data network, since in exascale systems, this factor determines energy consumption and influences different workload factors (POP, 2016)

Identifying the monitoring variables when running a process is the first step that must be performed to characterize the energy consumption processes according to the variables applied to the job (Slurms Guide, 2021). The first monitoring variable analyzed the use of nodes and the number of CPUs per node for each assigned job, finding that a given job, although it has an assigned number of CPUs, does not always use all of them, as shown in Fig. 2. To measure the use of nodes, we obtain the total number of CPUs assigned to a process, or the "total number of CPUs allocated to the job (AllocCPUs)" (Criado, et al., 2020), which is also referred to as NCPUs. To determine the number of nodes assigned to each job, we calculate the "number of nodes allocated to the job/step (AllocNodes)" (Slurms Guide, 2021).

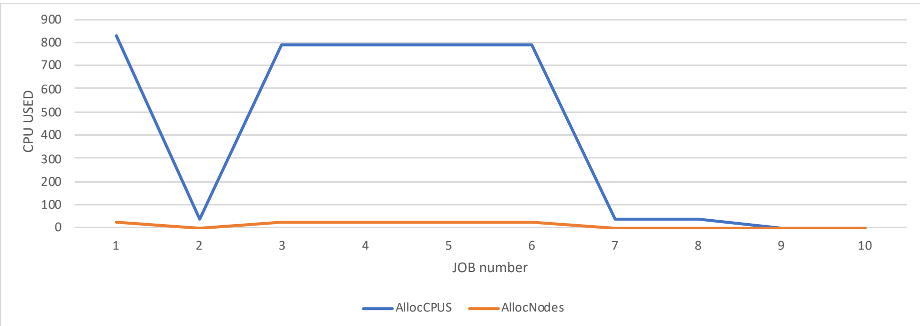

Figure 2. The number of CPUs and nodes used for each job.

Once the number of CPUs used for each job has been identified, we determine how long each CPU is used by the job that is assigned; this can vary according to the type of execution that is being carried out. To measure this, the CPU time (or process time) is defined as the amount of time for which a CPU was used to process the instructions of a computer program or operating system. For the data collection system, CPUTimeRAW is the time used (elapsed time * CPU count) by a job or step in CPU seconds, and DBIndex is a unique database index for entries in the job table (Carastan-Santos & Pham, 2022; Ficher, Berthoud, Ligozat, Sigonneau, & Wisslé, 2021). The execution time of each CPU for each job can be viewed as a set of processes to determine how these times affect the energy consumption; this can be seen in Fig. 3.

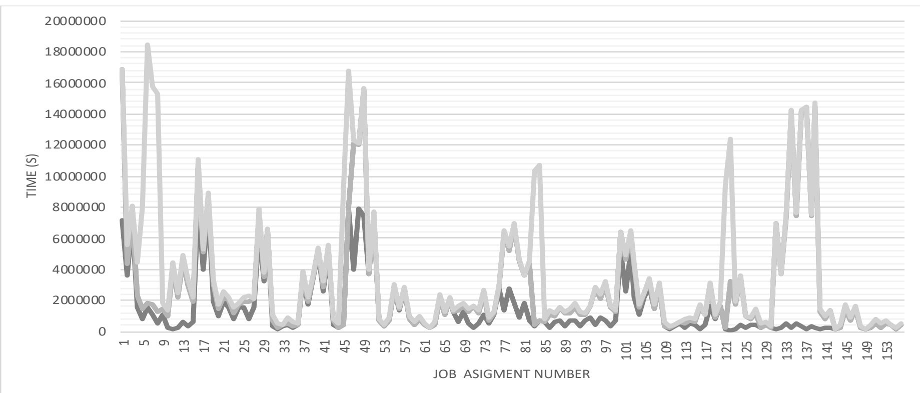

Figure 3. CPU time (raw usage) per month for each job.

The energy consumption related to the job being processed is analyzed by integrating the time of use and how much physical reading and writing it represents, and it can be visualized in terms of the energy consumption per CPU. This visualization makes it possible to analyze a set of factors for the energy consumption and calculate their correlations, as shown in Fig. 4. The node analysis involves projecting the energy consumption of all the nodes assigned to the job and all the CPUs that make up each node.

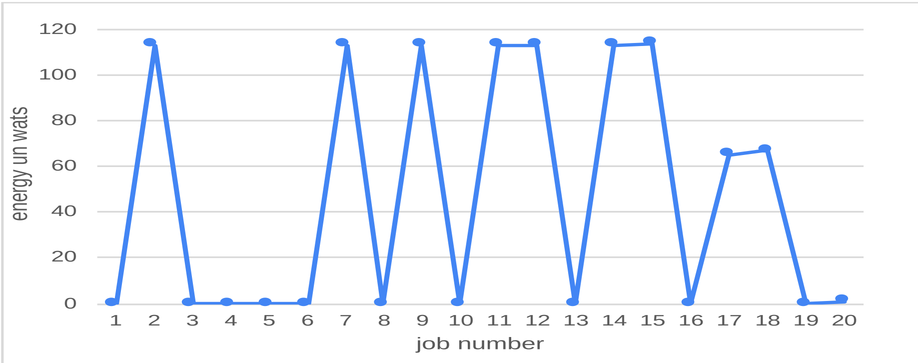

Figure 4. Energy consumed per node for each job.

The energy consumed by each job represents how much each node uses the CPUs assigned to it; this energy consumption is determined by the usage times of the CPUs. In this way, each job has a certain load usage time, and depending on the workload that is assigned to the node, the job consumption will be as the type of use of the assigned CPUs. In this way, it is possible to determine the usage time of the CPUs assigned to each node and associate it with the corresponding job, as shown in Fig. 5.

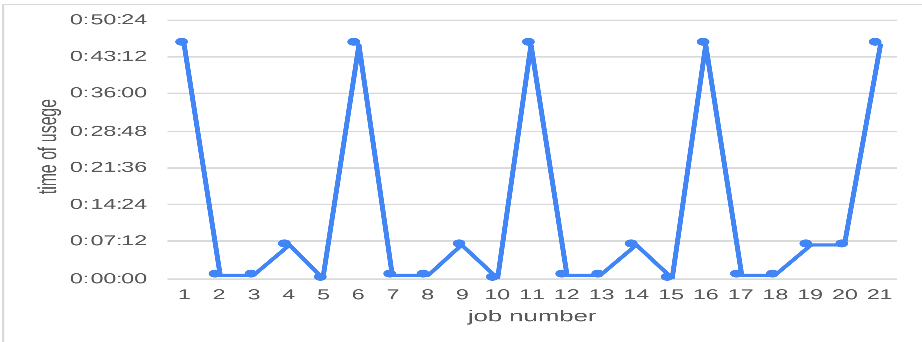

Figure 5. Usage times of nodes for each job.

To characterize the variables of the energy consumption during the execution of jobs, we can consider a series of activities carried out by processors that generate energy consumption; these factors determine the time of use and the type of process, which can help us break down what causes this energy consumption. In the case of disk reading, which is carried out to obtain data and send these data to be processed, and disk writing, which provides the generated data, the node is assigned certain hardware resources that are used for this function, although they are not in use for the entire time that they are assigned. Being an assigned resource and not being in use, this resource does not represent the same energy consumption. AveDiskRead and AveDiskWrite represent the average disk usage for reading data and the average disk usage for writing data, respectively. These values are defined for each assigned job and represent a series of functions, as shown in Fig. 6.

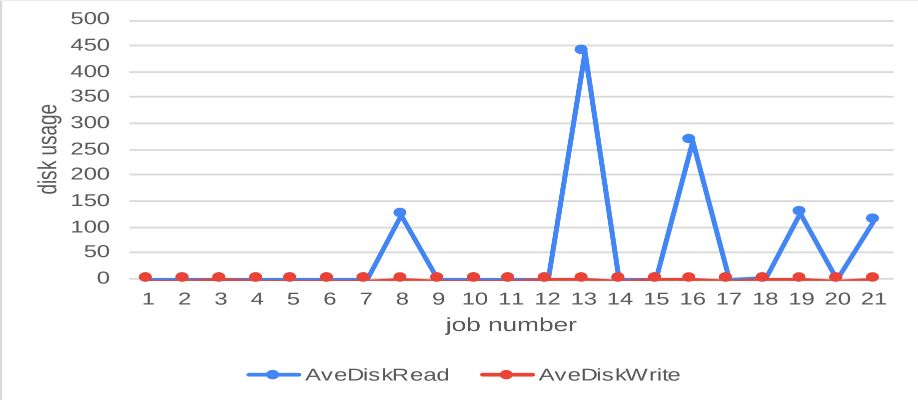

Figure 6. Disk read and disk write usage for each job.

This analysis considers a process that is composed of a series of repetitive processes based on processing iterations. This behavior represents a periodic request for the same number of resources with the same number of nodes and CPUs for a large number of assigned jobs. The behavior is generally very stable. The CPU and node usage are shown in Fig. 7. However, the disk read and write usage, as well as network and memory factors, vary because the allocated process contains different command schedules for processor usage.

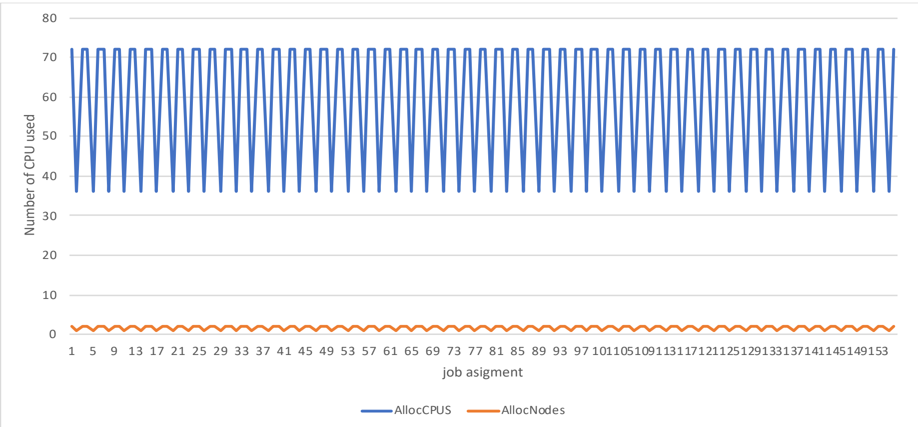

Figure 7. AllocCPUs and AllocNodes for each job in a synchronous process.

The same repetitive process involving the use of CPUs and assigned nodes that we analyze in Fig. 7 is made up of a series of reading and writing variables, as well as network use, which cause changes in the energy consumption. We see in Fig. 8 that all the processes, even those with the same CPU usage, do not represent the same energy consumption; here, differences are also identified in the architectures of the processors that are combined when the jobs are assigned.

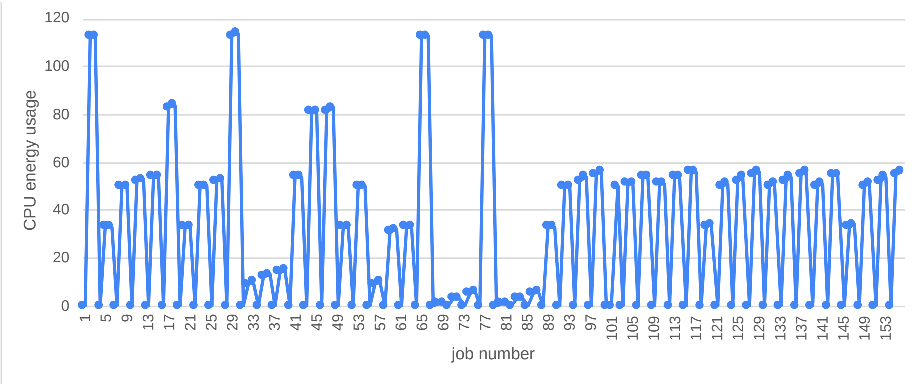

Figure 8. CPU energy usage for each job.

The total energy consumption of a job is made up of the factors mentioned in Fig. 8. We can see how the energy consumption of the CPUs for each job varies, and even though the allocation of resources is stable, the energy consumption is variable. This is due to the structure of the workload, the different software elements and the management of the operation of an HPC system. One can determine the energy consumption according to each type of job, considering not only how many CPUs have been assigned but also that the interference of the other processes means that there are significant variations in the energy consumption behavior, as shown for job 65 and job 84 in Fig. 8, which correspond to almost identical processing assignments, as shown in Fig. 7.

This energy consumption behavior can be controlled or predicted through the characterization of the energy consumption according to the software elements that make up the assigned workload, which is added to the number of CPUs that make up the assigned node. In this way, an analysis model based on hardware performance variables and a characterization of the functions that make up the workload can be used to describe and predict the average energy consumption per job. Doing this on a large scale in an HPC system would allow predictions to be made with different techniques, such as machine learning or deep learning techniques, to make the energy consumption more efficient in highly specialized architectures and with different processors and pieces of equipment that are in operation in the analyzed data center.

4 Conclusions

The different performance analysis techniques that are applied to HPC systems allow us to visualize how the assigned jobs are carried out, and in this way, we can generate energy consumption reports associated with each architecture and type of software implemented.

The behavior of the energy consumption in the analyzed HPC system is directly affected by the type of job that is processed; we can see this when we analyze the reading and writing variations of the processing variables in a system (AllocCPUs) and the number of nodes assigned to each job (AllocNodes).

Writing to the disk and reading from the disk generate an amount of energy consumption that is proportional to the number of assigned nodes; however, depending on the requested job and its concurrency, this energy consumption either stabilizes or increases.

The analysis of the energy consumption is influenced by the type of software that contains the workload; its characteristics directly affect the energy consumption of each node since the use of CPUs and reading and writing operations varies as the equipment processes the jobs. Controlling this type of behavior represents an assignment challenge that involves the development and scheduling of jobs; controlling these behaviors, as shown in Fig. 6, can make the energy consumption more stable, which can have a large impact on the total energy consumption of an HPC system.

5 References

Banchelli, F., Garcia-Gasulla, M., Houzeaux, G., & Mantovani, F. (2020). Benchmarking of state-of-the-art HPC Clusters with a Production CFD Code. *PASC '20: Proceedings of the Platform for Advanced Scientific Computing Conference. Article No. 3*, pp. 1-11. Geneva, Switzerland: Association for Computing Machinery. doi:10.1145/3394277.3401847

Carastan-Santos, D., & Pham, T. T. (2022). Understanding the Energy Consumption of HPC Scale Artificial Intelligence. In P. Navaux, C. J. Barrios H, C. Osthoff, & G. Guerrero (Ed.), *High Performance Computing. CARLA 2022. Communications in Computer and Information Science. 1660*, pp. 131-144. Springer, Cham. doi:10.1007/978-3-031-23821-5_10

Criado, J., Garcia-Gasulla, M., Kumbhar, P., Awile, O., Magkanaris, I., & Mantovani, F. (2020, September 14). CoreNEURON: Performance and Energy Efficiency Evaluation on Intel and Arm CPUs. *2020 IEEE International Conference on Cluster Computing (CLUSTER)* (pp. 540-548). Kobe, Japan: IEEE. doi:10.1109/CLUSTER49012.2020.00077

D'Agostino, D., Quarati, A., Clematis, A., Morganti, L., Corni, E., Giansanti, V., . . . Merelli, I. (2019). SoC-based computing infrastructures for scientific applications and commercial services: Performance and economic evaluations. *Future Generation Computer Systems, 96*, 11-22. doi:10.1016/j.future.2019.01.024

Dawson, W., Mohr, S., Ratcliff, L. E., Nakajima, T., & Genovese, L. (2020). Complexity Reduction in Density Functional Theory Calculations of Large Systems: System Partitioning and Fragment Embedding. *Journal of Chemical Theory and Computation, 16, 5*, 2952-2964. doi:10.1021/acs.jctc.9b01152

Djurfeldt, M., Johansson, C., Ekeberg, Ö., Rehn, M., Lundqvist, M., & Lansner, A. (2005). *Massively parallel simulation of brain-scale neuronal network models.* Report number: TRITA-NA-P0513, CBN, Royal Institute of Technology (KTH), Stockholm University, Computational Biology and Neurocomputing, School of Computer Science and Communication (CSC), Stockholm, Sweden. Retrieved from https://kth.diva-portal.org/smash/record.jsf?pid=diva2%3A220701&dswid=-2798

Favaro, F., Dufrechou, E., & Oliver, J. P. (2022). Time-Power-Energy Balance of BLAS Kernels in Modern FPGAS. In P. Navaux, C. J. Barrios H, C. Osthoff, & G. Guerrero (Ed.), *High Performance Computing. 9th Latin American Conference, CARLA 2022, Porto Alegre, Brazil, September 26–30, 2022, Revised Selected Papers. 1660*, pp. 78-89. Springer, Cham. doi:10.1007/978-3-031-23821-5_6

Ficher, M., Berthoud, F., Ligozat, A.-L., Sigonneau, P., & Wisslé, M. (2021). Assessing the carbon footprint of the data transmission on a backbone network. *2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)* (pp. 105-109). Paris, France: IEEE. doi:10.1109/ICIN51074.2021.9385551

Funika, W., Zientarski, M., Badia, R. M., Labarta, J., & Bubak, M. (2008). Performance Visualization Of Grid Applications Based On OCM-G And Paraver. In S. Gorlatch, P. Fragopoulou, & P. Thierry (Eds.), *Grid Computing* (pp. 109-120). Boston, MA, USA: Springer. doi:10.1007/978-0-387-09457-1_10

Hijji, M., Ahmad, B., Alwakeel, A., Alwakeel, M., Alharbi, L. A., Aljarf, A., & Khan, M. U. (2022). Cloud Servers: Resource Optimization Using Different Energy Saving Techniques. *Sensors, 22*(21), 8384. doi:10.3390/s22218384

Jarus, M., Varrette, S., Oleksiak, A., & Bouvry, P. (2013). Performance Evaluation and Energy Efficiency of High-Density HPC Platforms Based on Intel, AMD and ARM Processors. In J.-M. Pierson, G. Da Costa, & L. Dittmann (Ed.), *Energy Efficiency in Large Scale Distributed Systems. COST IC0804 European Conference, EE-LSDS 2013, Vienna, Austria, April 22-24, 2013, Revised Selected Papers. 8046*, pp. 182-200. Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-40517-4_16

Migliore, M., Cannia, C., Lytton, W. W., Markram, H., & Hines, M. L. (2006). Parallel Network Simulations with NEURON. *Journal of Computational Neuroscience, 21*, 119-129. doi:10.1007/s10827-006-7949-5

Mohr, S., Dawson, W., Wagner, M., Caliste, D., Nakajima, T., & Genovese, L. (2017). Efficient Computation of Sparse Matrix Functions for Large-Scale Electronic Structure Calculations: The

CheSS Library. *Journal of Chemical Theory and Computation, 13, 10*, 4684-4698. doi:10.1021/acs.jctc.7b00348

Plesser, H. E., Eppler, J. M., Morrison, A., Diesmann, M., & Gewaltig, M.-O. (2007). Efficient Parallel Simulation of Large-Scale Neuronal Networks on Clusters of Multiprocessor Computers. In A.-M. Kermarrec, L. Bougé, & T. Priol (Ed.), *Euro-Par 2007 Parallel Processing. Euro-Par 2007. Lecture Notes in Computer Science. 4641*, pp. 672-681. Springer, Berlin, Heidelberg. doi:10.1007/978-3-540-74466-5_71

POP. (2016, December 21). *Performance Optimisation and Productivity. A Centre of Excellence in HPC*. Retrieved June 21, 2023, from POP Homepage: https://pop-coe.eu/

Slurms Guide. (2021, June 29). Retrieved June 20, 2023, from https://slurm.schedmd.com/sacct.html

Universidad de Guadalajara. (2018, October 11). *CADS*. Retrieved June 20, 2023, from CADS Homepage: http://cads.cgti.udg.mx/

Wagner, M., Mohr, S., Giménez, J., & Labarta, J. (2019). A Structured Approach to Performance Analysis. In C. Niethammer, M. M. Resch, W. E. Nagel, H. Brunst, & H. Mix (Ed.), *Tools for High Performance Computing 2017. PTHPC 2017* (pp. 1-15). Dresden, Germany: Springer. doi:10.1007/978-3-030-11987-4_1